

Optimizing Service Placement in Edge-to-Cloud AR/VR Systems Using a Multi-Objective Genetic Algorithm

Mohammadsadeq G. Herabad¹^a, Javid Taheri^{1,2}^b, Bestoun S. Ahmed^{1,3}^c and Calin Curescu⁴^d

¹*Department of Mathematics and Computer Science, Karlstad University, Karlstad, Sweden*

²*School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Belfast, U.K.*

³*Department of Computer Science, Faculty of Electrical Engineering, Czech Technical University, Prague, Czech Republic*

⁴*Ericsson AB, Stockholm, Sweden*

Keywords: Edge-to-Cloud Computing, Service Placement, Multi-Objective Genetic Algorithm, Augmented Reality, Virtual Reality.

Abstract: Augmented Reality (AR) and Virtual Reality (VR) systems involve computationally intensive image processing algorithms that can burden end-devices with limited resources, leading to poor performance in providing low latency services. Edge-to-cloud computing overcomes the limitations of end-devices by offloading their computations to nearby edge devices or remote cloud servers. Although this proves to be sufficient for many applications, optimal placement of latency sensitive AR/VR services in edge-to-cloud infrastructures (to provide desirable service response times and reliability) remain a formidable challenging. To address this challenge, this paper develops a Multi-Objective Genetic Algorithm (MOGA) to optimize the placement of AR/VR-based services in multi-tier edge-to-cloud environments. The primary objective of the proposed MOGA is to minimize the response time of all running services, while maximizing the reliability of the underlying system from both software and hardware perspectives. To evaluate its performance, we mathematically modeled all components and developed a tailor-made simulator to assess its effectiveness on various scales. MOGA was compared with several heuristics to prove that intuitive solutions, which are usually assumed sufficient, are not efficient enough for the stated problem. The experimental results indicated that MOGA can significantly reduce the response time of deployed services by an average of 67% on different scales, compared to other heuristic methods. MOGA also ensures reliability of the 97% infrastructure (hardware) and 95% services (software).

1 INTRODUCTION

Augmented Reality (AR) and Virtual Reality (VR) technologies revolutionise our interaction with the digital and physical world. Integration of edge and cloud computing has led to significant progress in AR and VR systems, allowing low latency data processing, which is a crucial consideration in almost all AR/VR systems (Elawady and Sarhan, 2020).

Edge computing pushes data processing closer to its source, typically at the network edge, rather than relying on a centralized cloud server. On the network's edge, necessary computational tasks can

be executed on edge devices and/or nearby edge servers. While edge computing handles tasks requiring low latency, cloud computing can manage resource-intensive operations through powerful computational servers in data centers (Ji et al., 2023). Due to their complementary nature, integrating edge and cloud computing with AR/VR systems can address critical latency requirements, while enhancing AR/VR experiences by providing robust computational resources at both layers.

The complexities of real-time multimedia within AR/VR systems demand more careful consideration of how computational tasks should be distributed for execution to achieve minimum latency. This is an important concern to prevent disruption or delay in visual or interactive elements (Siriwardhana et al., 2021). Furthermore, infrastructure and service-related reliability play an important role in providing

^a  <https://orcid.org/0000-0002-2336-2077>

^b  <https://orcid.org/0000-0001-9194-010X>

^c  <https://orcid.org/0000-0001-9051-7609>

^d  <https://orcid.org/0000-0001-5129-2718>

a seamless and effective user experience (Dong et al., 2019; Khaleel, 2022). The hardware components of the edge-to-cloud infrastructure, from the input devices to the computing nodes and the output devices, must operate uninterruptedly to provide a smooth experience. Software running AR/VR services must also be highly reliable to prevent errors and corruptions that could disrupt the user experience. The system becomes even more complex due to heterogeneous devices and computing servers coupled with various communication capacities (Liu et al., 2023a).

To address the above challenges, adopting an optimization approach becomes essential to make optimal policies for placing AR/VR services on heterogeneous resources throughout the system. The optimization approach must determine decisions considering various factors including, but not limited to, memory and computing capacity of the nodes, communication characteristics between the nodes, as well as diverse service/resource (e.g., CPU, memory, disk, and network) requirements and characteristics related to required/running services. Therefore, making optimal decisions about the placement of services across available resources in the edge-to-cloud continuum can be extremely challenging and, therefore, identified (proved to be) as an NP-Hard problem (Raju and Mothku, 2023; Liu et al., 2023b; Huangpeng and Yahya, 2024).

Although the complexity of edge-to-cloud AR/VR systems can be studied for various applications and use cases, this paper, specifically, focuses on an AR/VR-assisted remote repair/maintenance use case in industrial sectors, such as those studied in (Fang et al., 2020). In our use case, developed as part of a collaborative project with Ericsson, when a malfunctioning device, within an industrial application, requires repair and no expert is available on-site, a nearby technician uses an AR/VR application on their device to connect with a remote expert. The technician shares enhanced videos of the malfunctioning device with the remote expert to solicit their help for identifying/troubleshooting the instrument, and/or developing a repair plan accordingly. In this specific use case, the latency of the AR/VR application plays a vital role in its functionality as a repairing tool, and thus the optimal placement of AR/VR services for efficient offloading of the computations over the network to meet service requirements (e.g., service level agreements) is essential.

This paper models the problem from a system design perspective and develops a Multi-Objective Genetic Algorithm (MOGA) to simultaneously address latency- and reliability-related concerns. The paper also demonstrates that heuristic algorithms are insuffi-

cient for addressing the complexities of service placement in edge-to-cloud AR/VR systems, thus justifying the use of meta-heuristic methods to achieve optimal or near-optimal solutions. We use MOGA because GA-based algorithms have been proven to be highly adaptable meta-heuristic algorithms that have been successfully applied to solve a diverse range of optimization problems. The capability of GA to handle large-scale problems and its inherent structure for parallelism also makes it a preferable choice for addressing the stated problem in this paper. Our contributions in this paper can be highlighted as follows:

- We modelled a multi-tier edge-to-cloud infrastructure with the focus on AR/VR systems that need to support service components with multiple versions (e.g., different codecs).
- We proposed a multi-objective GA-based service placement approach to simultaneously minimize both service response time and maximize infrastructure/service reliability.
- We proposed a grid-based fine-tuning procedure to identify optimal configurations (population size, crossover rate, etc.) for our MOGA, aiming to achieve a balance between the fitness and runtime of MOGA.
- We developed a tailor-made simulator to evaluate the performance of MOGA as compared with other heuristic algorithms. We use multiple metrics and explain how/why MOGA provides an optimal service placement solution in edge-to-cloud AR/VR systems.

The remainder of this paper is structured as follows. Section 2 provides the related work. Section 3 formulates the infrastructure and the service model. Section 4 defines the objective function. In Section 5, our solution is introduced. Section 6 outlines the experimental setup for simulations and evaluations, providing an analysis of the results obtained through the simulations. Finally, Section 7 presents the paper conclusion.

2 RELATED WORK

Over the past few years, many approaches were designed to use different techniques for offloading video streams from end-devices to the near edge devices or clouds to reduce service delays and service response time in edge-to-cloud environments (Ren et al., 2021; Acheampong et al., 2023; Chen et al., 2022; Huang and Friderikos, 2021). Specifically, the authors of (Cozzolino et al., 2022) proposed an approach called

Nimbus, which is a multi-objective solution for task allocation in edge-to-cloud environments. The primary goal of Nimbus is to minimize the latency of AR applications by offloading computational tasks to edge or cloud servers. This work, however, focuses on simple AR tasks without paying sufficient attention to the complications associated with multi-version AR tasks and procedures. Additionally, authors of (Yeganeh et al., 2023) devised an approach that uses Q-learning to reduce task execution time and mitigate energy consumption on end-devices by offloading computational tasks to edge-nodes. Through a series of experiments, the authors claimed that their approach provides a better solution, as compared with similar algorithms that were implemented in the article in 90% of the cases. The paper addresses the problem from a general perspective, without explicitly focusing on the complexity of video streaming.

Authors of (Mahjoubi et al., 2022b) formulated service placement as a Mixed-Integer Linear Programming (MILP) problem in edge-to-cloud computing. Subsequently, a single-objective genetic algorithm (Mahjoubi et al., 2022b) and a simulated annealing algorithm (Mahjoubi et al., 2022a) were used to solve the problem. These works focused on single-objective algorithms, aiming to minimize the runtime of the system without addressing other metrics, such as system reliability, in terms of software and hardware aspects. The authors of (De Souza et al., 2023) also introduced an approach based on Bee Colony optimization to reduce the execution time of the application for computational tasks by offloading them to the edge. The extensive experiments described in the paper supported the claim that their technique can reduce the execution time by 56% when compared with other heuristics. This work neglected to take into account the computational capacity of the cloud, particularly in scenarios with substantial computational demands. This is because the algorithm primarily focuses on the resource conditions of end devices and edge servers to make offloading decisions. The authors of (Fan et al., 2022) presented a collaborative approach based on the Lyapunov optimization technique for the placement of services within edge-to-cloud systems. Although this work demonstrated the effectiveness of the proposed algorithm in various scenarios to minimize overall task processing delay and ensure long-term task queue stability, it focused on single-objective scenarios where the main objective is to reduce the system delay. Similarly, the authors of (Wang et al., 2023) formulated service placement as a linear integer programming problem and proposed a polynomial-time method to make decisions collaboratively among edge nodes and handle

service heterogeneity in the system. This research focuses on service placement in small-scale scenarios, neglecting to explicitly consider aspects related to software and hardware reliability.

Besides approaches for single-objective problems, there were also several works focusing on multi-objective ones, such as (Abedi et al., 2022; Li et al., 2022; Lavanya et al., 2020; Madni et al., 2019). Despite the significant body of research investigating service placement and offloading strategies within edge-to-cloud computing environments, a limited portion of these studies consider the complexities of edge-to-cloud AR/VR systems. In fact, many studies adopted a broad perspective, failing to consider the unique challenges and requirements posed by the nature of AR/VR workloads. As a result, there is a notable gap in the literature on approaches that address the specific characteristics of AR/VR systems in terms of service response time, hardware reliability, and service reliability, specifically in cases where multiple versions of AR/VR service components (co-)exists in a platform. This gap deserves further investigation and exploration in this domain, and motivated us to design algorithms in this article.

3 SYSTEM MODEL

This section presents mathematical models for a multi-tier infrastructure that hosts AR/VR-based services.

3.1 User and Helper Model

We consider a use case where ‘users’ can connect with remote ‘helpers’ using AR/VR applications. This connection allows users to consult with remote helpers (called user-helper pairs) to perform repair or maintenance tasks on industrial sites.

Both users and helpers have their own devices, denoted by $U = \{u_1, u_2, \dots, u_n, \dots, u_N\}$ and $H = \{h_1, h_2, \dots, h_m, \dots, h_M\}$, respectively. Resources for users and helpers are modeled as nodes with computation capacity (CC), memory capacity (MM) and disk capacity (DC), as well as other characteristics such as an Operating System (OS) and a reliability score (RS). These characteristics are indicated by $u_n = \langle CC_n, MC_n, DC_n, OS_n, RS_n \rangle$ and $h_m = \langle CC_m, MC_m, DC_m, OS_m, RS_m \rangle$; we define U^{ch} and H^{ch} (Equations (1) and (2)) to represent the characteristics of all users and helpers, respectively. We also define $P = \{p_1, p_2, \dots, p_i, \dots, p_I\}$ to show the set of pairs, where $p_i = \langle u_n, h_m \rangle$ represents the i^{th} pair. Table 1 describes the notation related to the user and helper

Table 1: Notations related to the user and helper nodes.

| Notation | Description |
|----------|---|
| U | The set of user nodes |
| H | The set of helper nodes |
| u_n | The n^{th} user node |
| h_m | The m^{th} helper node |
| U^{ch} | The characteristics of all user nodes |
| H^{ch} | The characteristics of all helper nodes |
| P | The set of pairs |
| p_i | The i^{th} pair |
| N | The total number of user nodes |
| M | The total number of helper nodes |
| I | The total number of pairs |
| CC | Computation capacity (MIPS) |
| MC | Memory capacity (GB) |
| DC | Disk capacity (GB) |
| RS | Reliability score |
| OS | Operating system |

nodes.

$$U^{ch} = \begin{bmatrix} CC_1 & MC_1 & DC_1 & OS_1 & RS_1 \\ CC_2 & MC_2 & DC_2 & OS_2 & RS_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ CC_N & MC_N & DC_N & OS_N & RS_N \end{bmatrix} \quad (1)$$

$$H^{ch} = \begin{bmatrix} CC_1 & MC_1 & DC_1 & OS_1 & RS_1 \\ CC_2 & MC_2 & DC_2 & OS_2 & RS_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ CC_M & MC_M & DC_M & OS_M & RS_M \end{bmatrix} \quad (2)$$

3.2 Service Model

The system hosts multiple AR/VR services, $S = \{S_1, S_2, \dots, S_x, \dots, S_X\}$, where S_x denotes the service running through the x^{th} user-helper pair. Each service, consists of Y service components, where SC_y^x represents service component y of x^{th} user-helper pair, $S_x = \{SC_1^x, SC_2^x, \dots, SC_y^x, \dots, SC_Y^x\}$. Similarly, every service component (in a service) has V versions, where $SC_{y,v}^x$ represents version v of the service component y of the service belonging to pair x , $SC_y^x = \{SC_{y,1}^x, SC_{y,2}^x, \dots, SC_{y,v}^x, \dots, SC_{y,V}^x\}$. Each version of a service component has specific requirements (demands) regarding computational power, memory capacity, and disk space to ensure successful execution; it also produces/transfers data (with specific

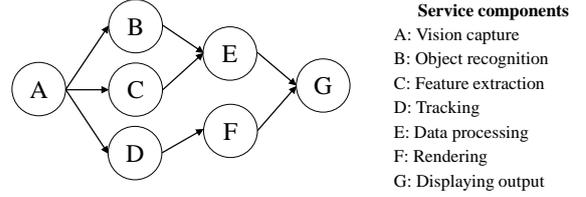


Figure 1: An example DAG for an AR/VR application.

size) to other dependent service components. Furthermore, each version is associated with various providers, including external providers such as AWS and Azure, and internal providers such as hosted Kubernetes clusters. There are also attributes for service components, such as Transcode, Codec Type, and a reliability score to indicate how they work and how often they fail. We define SC_{ych}^x to represent all the resource requirements and characteristics of all versions of a service component (Table 2 shows the notation of the service model.)

$$SC_{ych}^x = \begin{bmatrix} CR_1 & MR_1 & DR_1 & DS_1 & PR_1 & TC_1 & CT_1 & RS_1 \\ CR_2 & MR_2 & DR_2 & DS_2 & PR_2 & TC_2 & CT_2 & RS_2 \\ \vdots & \vdots \\ CR_V & MR_V & DR_V & DS_V & PR_V & TC_V & CT_V & RS_V \end{bmatrix} \quad (3)$$

Where S_{ych} denotes the resource requirements and characteristics of all service components belonging to the x^{th} user-helper pair.

$$S_{ych} = \begin{bmatrix} SC_{1ch}^x & SC_{2ch}^x & \dots & SC_{ych}^x & \dots & SC_{Ych}^x \end{bmatrix} \quad (4)$$

The service components of each service are interdependent. Therefore, each service can be represented by a Directed Acyclic Graph (DAG), where the vertices in the DAG represent the service components, and the links between the vertices represent the dependencies between the service components. Figure 1 shows an example of a DAG that represents relationships between the service components of an AR/VR application. In this paper, an upper binary triangular matrix is used to model the DAG, where a value of '1' in the matrix implies a connection between the corresponding service components.

3.3 Infrastructure and Network Model

We employ a three-tier infrastructure, including access points (AP), edge, and cloud tiers. Tier-1 is defined as a group of APs with specific computing capabilities. These APs establish direct wireless connections with users, acting as network entry points. Tier-2 comprises a collection of computing nodes positioned close to (in terms of network accessibility and speed) backbone routers in proximity to

Table 2: Notations related to service model.

| Notation | Description |
|-----------------|---|
| S | The set of services |
| S_x | The service belongs to the x^{th} user-helper pair |
| $S_{x^{ch}}$ | The characteristics of all service components of S_x |
| SC_y^x | The y^{th} service component of S_x |
| $SC_{y,v}^x$ | The v^{th} version of the y^{th} service component of S_x |
| $SC_{y^{ch}}^x$ | The characteristics of all versions of SC_y^x |
| X | The total number of services ($X = I$) |
| Y | The total number of service components of each service |
| V | The total number of service component versions |
| CR | Computation requirement of a service component (MIPS) |
| MR | Memory requirement of a service component (GB) |
| DR | Disk requirement of a service component (GB) |
| DS | The data size that is required to transfer (Mb) |
| PR | Service component provider |
| TC | Service component transcode |
| CT | Service component codec type |

Tier-1. Nodes within Tier-2 have enhanced computational and storage capacities. Tier-3 contains the cloud in our system, providing the highest computational power and storage capacity among all tiers. The increase in computing capability in the upper tiers comes with a rise in data transmission latency. We assume that the helpers are connected (accessible) through the cloud (because they are located in remote places). Figure 2 illustrates the entities in our system.

The computing nodes in the different tiers have their own device characteristics, denoted as $CN = \{CN_1, CN_2, \dots, CN_k, \dots, CN_K\}$, where $CN_k = \langle CC_k, MC_k, DC_k, OS_k, RS_k \rangle$. Moving to the upper tiers, the memory and computational capacity of computing nodes increase with respect to CPUs, GPUs, and RAM, while their network accessibility decreases (i.e., communicating with them incurs more delay). We define CN^{ch} to represent the characteristics of the computing nodes within the infrastructure.

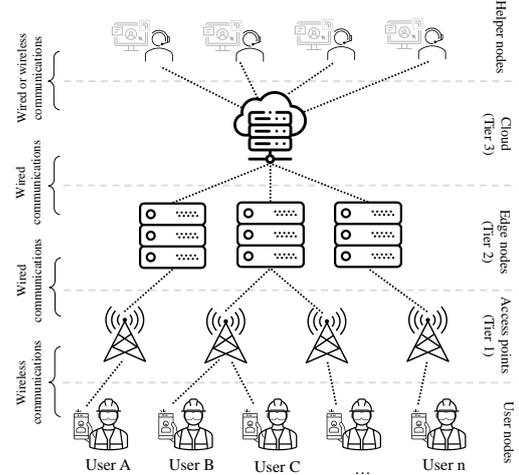


Figure 2: Multi-tier edge-to-cloud infrastructure.

$$CN^{ch} = \begin{bmatrix} CC_1 & MC_1 & DC_1 & OS_1 & RS_1 \\ CC_2 & MC_2 & DC_2 & OS_2 & RS_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ CC_K & MC_K & DC_K & OS_K & RS_K \end{bmatrix} \quad (5)$$

To model the network communications among computing nodes within the infrastructure and between users, helpers, and computing nodes, we define CN^L (Equation 6) to represent the available bandwidth (BW) and transmission delays (LD) on each communication link. For example, $BW_{2,1}$ shows the available bandwidth between computing nodes 2 and 1, similarly, $LD_{2,1}$ shows observed delay between computing nodes 2 and 1. In Equation 6 that represents CN^L , we assume that rows and columns from 1 to K are related to computing nodes, columns from $K+1$ to $K+N$ are associated with user nodes, and columns from $K+N+1$ to $K+N+M$ are related to helper nodes. Here, each element represents available bandwidth and observed delay (e.g., half of the ICMP round-trip time between two entities) among users, helpers, and computing nodes. We use ICMP round-trip time to consider all potential delays between two nodes, including those resulting from routing processes. Since the users are not connected to each other and the helpers also are not connected to each other, the number of rows in the matrix is equal to K . Table 3 describes the notation related to the infrastructure and the network model.

$$CN^L = \begin{bmatrix} \langle 0, 0 \rangle & \cdots & \langle BW_{1,K+N+M}, LD_{1,K+N+M} \rangle \\ \langle BW_{2,1}, LD_{2,1} \rangle & \cdots & \langle BW_{2,K+N+M}, LD_{2,K+N+M} \rangle \\ \vdots & \ddots & \vdots \\ \langle BW_{K,1}, LD_{K,1} \rangle & \cdots & \langle BW_{K,K+N+M}, LD_{K,K+N+M} \rangle \end{bmatrix} \quad (6)$$

Table 3: Notations of the infrastructure and network model.

| Notation | Description |
|------------|--|
| CN | The set of computing nodes in the infrastructure |
| CN^{ch} | The resource characteristics of the computing nodes |
| CN^L | The characteristics of communication links among users, helpers, and computing nodes |
| $BW_{K,1}$ | The available bandwidth between node K and node 1 (Mbps) |
| $LD_{K,1}$ | The link delay between node K and node 1 (ms) |
| K | The number of total computing nodes |

4 OPTIMIZATION MODEL

We adopt a centralized service placement approach, where service placement algorithms operate in the cloud to decide on the placement of AR/VR services across the entire infrastructure. We assume that data related to services and infrastructure characteristics are available before making service placement decisions. To formulate an objective function for optimal decision-making and service placement, various factors are considered, and discussed in the following subsections.

4.1 Data Transmission Delay Model

The data transmission delay depends on both the network bandwidth and the data size of the service components. In addition, other factors, such as jitter, can affect the transmission time of data. Therefore, we consider the links' bandwidth, the data size of the service components, and the Round Trip Time (RTT), which can be measured proactively using ICMP messages, of the links to calculate transmission delays (Cozzolino et al., 2022). Equation (7) calculates the delay of data transmission, where DS_v shows the required data size transferred by version v of the service component and BW shows the bandwidth of the transmission link. Also, we divide the RTT by 2 to account for the one-way delay of the link.

$$TD(SC_{y,v}^x) = \frac{DS_v}{BW} + \frac{RTT}{2} \quad (7)$$

If a service component needs to transfer data to multiple service components in different places, equation (7) is calculated separately for each. The notations related to the transmission delay model are presented in Table 4.

Table 4: Notations of the transmission delay model.

| Notation | Description |
|------------------|--|
| DS_v | The data size transferred by v^{th} version of the service component |
| $TD(SC_{y,v}^x)$ | The transmission delay of $SC_{y,v}^x$ |
| BW | The link bandwidth |
| RTT | Round trip time of the link |

Table 5: Notations related to the execution time model.

| Notation | Description |
|------------------|--|
| CC_k | The computation capacity of computing node k |
| CR_v | The required computation of the v^{th} version of SC_y^x |
| W_v | The waiting time of the v^{th} version of SC_y^x |
| $ET(SC_{y,v}^x)$ | The execution time of $SC_{y,v}^x$ |

4.2 Execution Time Model

The execution time of a service component is closely related to the computational capacity of the computing node and the computational requirements of the service component. Consequently, the execution time of a service component is calculated using Equation (8). CR_v reflects the total number of instructions required to execute a version v of the service component, and CC_k reflects the computational capacity of the computing node k . We assume the computing nodes support multi-threading and can execute the service components concurrently. However, a waiting time is considered in execution time, denoted as W_v . This waiting time occurs when a service component has been placed on the computing node but has not started its execution on the node yet. The notations related to the execution time model are described in Table 5.

$$ET(SC_{y,v}^x) = \frac{CR_v}{CC_k} + W_v \quad (8)$$

4.3 Response Time Model

To calculate the response time of a service component, various delays are considered: provider delays and coding delays, denoted as $PD(SC_{y,v}^x)$ and $CD(SC_{y,v}^x)$, respectively. Provider delays are assumed to have (almost) constant values and can be estimated through passive measurements on the network. Coding delays reflect the execution time of specific versions of encoding and decoding algorithms; they also

Table 6: Notations related to the response time model.

| Notation | Description |
|------------------|---|
| $RT(SC_{y,v}^x)$ | The response time of $SC_{y,v}^x$ |
| $RT(S_x)$ | The response time of S_x |
| $RT(S)$ | The response time of S |
| $PD(SC_{y,v}^x)$ | The provider delay of $SC_{y,v}^x$ |
| $CD(SC_{y,v}^x)$ | The encoding and decoding delay of $SC_{y,v}^x$ |

assumed constant and can be obtained from the algorithm providers and/or through multiple runs on a sample video feed.

The response time for a service component is calculated using Equation (9). Similarly, the total response time for a service and all services are obtained by (10) and (11), respectively. Table 6 describes notations of the response time model.

$$RT(SC_{y,v}^x) = TD(SC_{y,v}^x) + ET(SC_{y,v}^x) + PD(SC_{y,v}^x) + CD(SC_{y,v}^x) \quad (9)$$

$$RT(S_x) = \sum_{y=0, v \in V}^Y RT(SC_{y,v}^x) \quad (10)$$

$$RT(S) = \sum_{x=0}^X RT(S_x) \quad (11)$$

4.4 Reliability Model

The reliability of both the nodes and the service components is measured by their historical performance. In particular, researchers of (Moghaddas et al., 2016; Amini Motlagh et al., 2022; Liu et al., 2015) have introduced methods to measure their reliability scores. Taking into account the interdependent relationships between the components of the service, the overall reliability of the service (software) is calculated using Equation (12) (Maciel et al., 2021) to determine the probability of successfully completing a service.

$$RS(S_x) = \prod_{y=0, v \in V}^Y RS_{SC_{y,v}^x} \quad (12)$$

The average reliability of all services is calculated using Equation (13).

$$RS(S) = \frac{\sum_{x=0}^X RS(S_x)}{X} \quad (13)$$

Taking into account the independent relationship among the computing nodes in the infrastructure, the total reliability of the computing nodes is represented

Table 7: Notations related to the reliability model.

| Notation | Description |
|-------------------|---|
| $RS_{SC_{y,v}^x}$ | The reliability score of $SC_{y,v}^x$ |
| RS_k | The reliability score of a computing node k |
| $RS(S_x)$ | The reliability score of S_x |
| $RS(S)$ | The average reliability of all services |
| $RS(CN)$ | The reliability of computing nodes |
| $RS(U)$ | The average reliability score of the users |
| $RS(H)$ | The average reliability score of helpers |
| $RS(P)$ | The total hardware reliability of the pairs |

by Equation (14) (Maciel et al., 2021), where RS_k measures the reliability score of a computing node k .

$$RS(CN) = 1 - (1 - RS_1) \times \dots \times (1 - RS_k) \times \dots \times (1 - RS_K) \quad (14)$$

In a user-helper pair, because the failure of one part (i.e., user node, helper node or computing nodes) can adversely affect other parts, the overall hardware reliability of the pairs is determined using Equation (15) (Maciel et al., 2021). Table 7 describes the notations of the reliability model.

$$RS(P) = RS(CN) \cdot RS(U) \cdot RS(H) \quad (15)$$

4.5 Objective Function

Our main objective is to minimize the total response time of all services, while maximizing both hardware and software reliability. Therefore, our primary function is to map service components to heterogeneous computing nodes, as formulated in Equation (16).

$$\text{Objective function} : \begin{cases} \min RT(S) \\ \max RS(S) \\ \max RS(P) \end{cases} \quad (16)$$

subjected to:

$$\sum_{y,v \in Y,V} SC_{y,v}^x = 1, \forall y \quad (17)$$

$$\sum_{x,y,v \in X,Y,V} res(SC_{y,v}^x) < res(CN_k), \forall k \quad (18)$$

$$SC_{y,v}^x = u_x, \forall x \quad (19)$$

$$SC_{y,v}^x = h_x, \forall x \quad (20)$$

The constraint (17) shows that a service component within the system must be exclusively assigned to a single computing node (no additional copies of a service component are allowed). The constraint (18) reflects that the total resources required

Table 8: Notations related to MOGA parameters.

| Notation | Description |
|-----------------|---|
| ps | Populations size |
| cr | Crossover rate |
| mr | Mutation rate |
| ss | Tournament selection size |
| it | Number of iterations |
| $f(i)$ | Fitness of i^{th} chromosome |
| w_1, w_2, w_3 | A value between 0 and 1 where $w_1 + w_2 + w_3 = 1$ |

for all service components running on the computing node ($\sum res(SC_{y,v}^x)$) must not exceed the available resources of the computing node ($res(CN_k)$). The constraint (19) shows that each user device is restricted to execute only service components that are directly associated with that user. Similarly, the constraint (20) explicitly requires helper devices to run service components that belong to that helper node and its user. Because the formalized objective function belongs to the class of problems (NP-Complete) that cannot be solved in polynomial time, we developed our approach (MOGA) to efficiently solve this problem.

5 PROPOSED SOLUTION (MOGA)

Our solution begins the optimization process by defining the objective function and generating the initial population. Each individual in the population is assessed by a fitness function. Genetic operators, including crossover and mutation, generate offspring solutions. The population is then updated by replacing individuals with new offspring using a selection operator. MOGA also applies a healing operator to individuals to heal (rather than discard) chromosomes that partially violate constraints. The algorithm continues to iterate via these steps until a termination criterion is fulfilled. Table 8 describes MOGA parameters.

5.1 Chromosome Encoding

Each chromosome is represented using an array of $X \times Y$ elements. Each element is considered as a gene, consisting of a tuple with two values: the first value indicates the versions of the service components and the second value is associated with the computing node running the version of the service component.

Figure 3 represents an example chromosome, where the element $x \times Y^{th}$ shows the Y^{th} service component of service x is using the version 3 and runs on the computing node 8.

To generate the initial population of chromosomes, random versions of service components are arbitrarily assigned to computing nodes. By introducing randomness in both the assignment of the versions and computing nodes, the initial population covers a broad search space, leading to the exploration of a diverse solution space.

5.2 Fitness Function

For each chromosome, three values are calculated: response time, infrastructure reliability, and service reliability (discussed in section 4). A lower response time and higher values for both infrastructure and service reliabilities indicate a superior chromosome. To this end, we transform the reliabilities' values to create a minimization fitness function. It is important to note that the reliabilities are constrained within the range of 0 to 1 ($1 \geq RS(P) > 0, 1 \geq RS(S) > 0$).

After calculating the values, normalization is performed to ensure consistent scaling. Subsequently, a fitness function is defined using the weighted sum approach, denoted in Equation (21), where i reflects the i^{th} chromosome and $f(i)$ is a combined fitness function of the chromosome. Given the importance of both response time and system reliability in improving overall service quality, we consider equal weights for $RT(S)$, $RS(P)$, and $RS(S)$, where $w_1 = 0.33, w_2 = 0.33$ and $w_3 = 0.33$ (although other weights can be set to match other use cases). Because the fitness function is a minimization function, MOGA operates to minimize its value, and thus lower fitness values indicate better chromosomes.

$$f(i) = w_1 \cdot RT(S) + w_2 \cdot (1 - RS(P)) + w_3 \cdot (1 - RS(S)) \quad (21)$$

5.3 Selection Operator

A tournament strategy is empirically chosen as the selection operator; Other selection procedures, such as the roulette wheel, led to the generation of more populations for convergence. In tournament selection, a number of chromosomes (ss) are randomly chosen and then ranked based on their fitness values. The one with the best fitness value is selected to be part of the next-generation population. These steps are repeated iteratively until the target population size (ps) is achieved for the next generation.

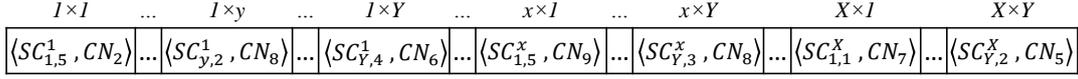


Figure 3: Chromosome encoding.

5.4 Crossover and Mutation Operators

Crossover operator can be accomplished using various techniques such as single-point crossover, multi-point crossover, and uniform crossover. In this paper, we empirically employed the single-point crossover technique, where two chromosomes (called parent pair) among the population with a crossover rate (cr) are selected for the crossover operation. The crossover point is randomly generated for the parent chromosomes. Genes located after the crossover point on one parent chromosome are exchanged with genes located before the crossover point in the other parent chromosome to generate two offspring. Then, the parent pair is replaced with the offspring in the population. We use the single-point crossover because it can make a trade-off between exploitation and the operator's run-time, particularly when dealing with substantial population sizes and extensive search spaces, making it a preferable choice over other crossover techniques. Other methods such as multi-point and uniform crossover also led to more complicated healing processes and consequently more convergence time.

The mutation operator randomly modifies the values of a gene, allowing the algorithm to explore a broader search space and preventing premature convergence to local optimum points. In MOGA, each gene on a chromosome is subjected to mutation with a mutation rate (mr). We utilize an insertion mutation technique that involves assigning random valid values to the selected gene.

5.5 Healing Operator

After crossover and mutation operators, it is inevitable that some of the offspring may not satisfy all the constraints of the objective function. Therefore, in this paper, a healing operator is designed to guarantee the satisfaction of all constraints by modifying some genes on such chromosomes. Since the constraints (17) are met during the chromosome encoding process, the healing operator is designed to satisfy the constraints (18), (19), and (20). Specifically, the healing operator examines the chromosome after applying the crossover and mutation operators. If constraint (18) is not met for a specific user/helper/computing node, service components of the node is reassigned to other nodes with available computing capabilities

until constraint (18) is satisfied. Then it updates the chromosome based on the new assignment. Similarly, if a service component does not belong to a user-helper pair but is running on the user or helper node, it is assigned to other nodes to satisfy the constraints 19 and 20. After the healing operator, the new generation of chromosomes is determined through the selection operator.

6 EXPERIMENTAL SETUP

For the evaluation, we implemented a tailor-made simulator in Node.js to simulate the entire infrastructure and services. The simulator was implemented to precisely match our problem, reflecting the desired scenarios and objectives. The simulator accepts inputs in the form of JSON objects/files where the infrastructure, the AR/VR service, and the users/helpers are defined. We designed a client-server architecture for the simulator, where the client-side sends all information of a problem instance (the infrastructure, service characteristics, etc.) to the server-side in the JSON format. Subsequently, the server employs various solvers, including MOGA, to determine the optimal solution. The results are then sent back to the client side. The simulator is containerized and, along with all problem instances in this paper, are made available for access from a GitHub (Herabad, 2024) address. WiKis and YAML files are also provided to facilitate deploying the prepared docker-images on Kubernetes platforms.

6.1 Other Scheduling Algorithms

To gauge the efficiency of MOGA and compare its results, we also designed and implemented other heuristic solvers; namely, (1) Task Continuation Affinity (TCA), (2) Least Required CPU (LRC), (3) Most Data Size (MDS), (4) Most Reliability (MR), (5) Most Powerful (MP), and (6) Least Powerful (LP). TCA executes the first version of the service component on the user nodes if sufficient resources are available. In the cases of inadequate resources, it attempts with subsequent versions. If there are still insufficient resources after checking all versions, the service component is placed to the upper tier. Similarly, if the computing nodes in each tier do not have sufficient resources after checking all versions, the remaining

Table 9: Evaluation scenarios and implementation setup.

| Specifications | Small-scale | Medium-scale | Large-scale | xLarge-scale |
|-------------------------|---|--------------|-------------|--------------|
| N/M | 15/8 | 50/25 | 100/50 | 250/125 |
| No. of CN in Tier-1/2/3 | 10/8/2 | 30/18/4 | 75/60/8 | 150/100/15 |
| X/Y | 15/5 | 50/5 | 200/5 | 250/5 |
| V | 5 | 6 | 7 | 8 |
| CR of SC | [800 - 3000] (MIPS) | | | |
| MR of SC | [1.5 - 3.3] (GB) | | | |
| DR of SC | [1 - 3] (GB) | | | |
| CC of U/H | [500 - 2200]/[1500 - 2500] (MIPS) | | | |
| MC of U/H | [2 - 4]/[2 - 4] (GB) | | | |
| DC of U/H | [4 - 8]/[4 - 8] (GB) | | | |
| CC of CN in Tier-1/2/3 | [1500 - 2000]/[5000 - 15000]/[15000 - 30000] (MIPS) | | | |
| MC of CN in Tier-1/2/3 | [4 - 8]/[8 - 16]/[32 - 64] (GB) | | | |
| DC of CN in Tier-1/2/3 | [8 - 32]/[32 - 128]/[128 - 256] (GB) | | | |
| DS of SC | [500 - 800] (Mb) | | | |
| RS of CN/SC | [0.7 - 0.9]/[0.9 - 0.99] | | | |
| Link BW | [100 - 500] (Mbps) | | | |
| Link RTT | [500 - 1200] (ms) | | | |

service components are placed to the upper tier. LRC operates similarly to TCA, but only selects the version of service components that demand the least CPU for completion. MDS also operates similarly to TCA, but prioritizes running service components with the largest data size on user nodes or closer tiers to the users as long as resources are available. MR runs the version of the service components with the highest reliability on the computing node that has the highest reliability in the infrastructure. MP runs the least computationally demanding version on the most powerful node in terms of computational capacity. LP is the opposite of MP, where the most computationally demanding version is executed on the least powerful node.

6.2 Evaluation Scenarios

To evaluate the performance of MOGA, we consider four evaluation scenarios (i.e., problem instances): small-scale, medium-scale, large-scale, and xLarge-scale scenarios. Table 9 shows the specifications of all scenarios along with the characteristics of their assumed infrastructure and services. The values for the infrastructure reflect the scale (properties of all nodes and networking links), and thus selected carefully (based on our discussion with our industry partners) to reflect realistic infrastructures. To pressure all

algorithms and push them to their limits (in terms of the quality of the solutions and the convergence time), our simulator then generates the service components and their requirements so that the total resource capacity of the service components constitutes 60-70% of the total resource capacity of the infrastructure. Additionally, we assume that each helper can provide assistance to more than one user; that is, a helper can be part of several user-helper pairs.

6.3 Configuration Setup of MOGA

The configuration of MOGA significantly affects its performance. We designed a grid-based tuning strategy to find the best configuration for our MOGA. To this end, we first ran the grid-based tuning for all scenarios and investigated the impacts of the parameters on the performance of MOGA in each scenario.

In MOGA, similar to any other GA-based solver, the fitness value tends to improve as the population size increases, and larger population sizes (almost) always yield better solutions. However, running MOGA for large and extra-large (xLarge) scenarios with large population sizes led to unacceptable (more than 24 hours) convergence times. Our grid-based tuning procedure is to strike a balance between MOGA's quality of solutions and its convergence time. This is because some solutions (fitness points)

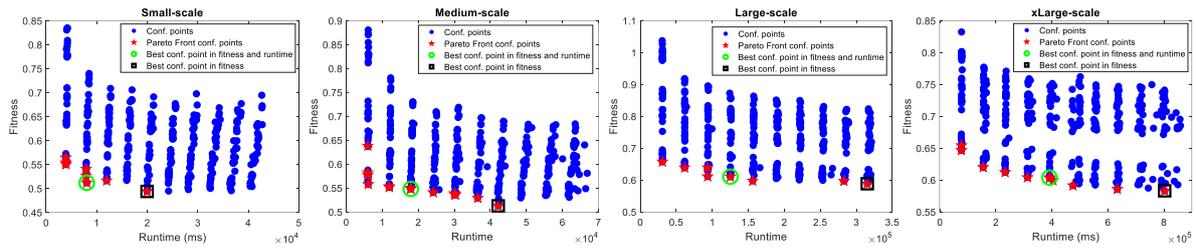


Figure 4: Selecting the best configuration for MOGA in different scales by the Pareto Front-based method.

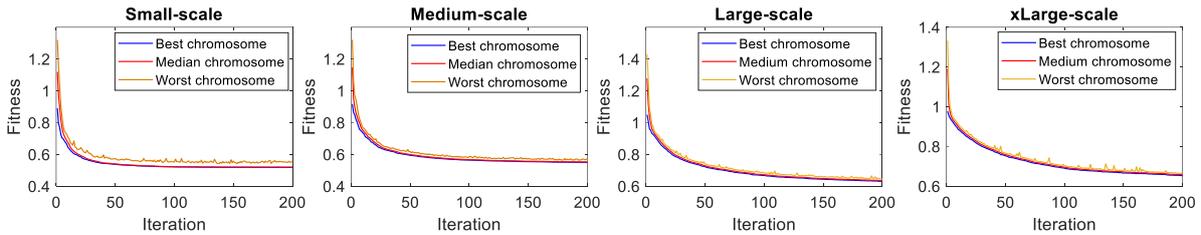


Figure 5: Convergence process of MOGA in different scales.

produced through smaller population sizes may not differ much from better solutions found through larger population sizes at the cost of much higher convergence time. This implies that acceptable (not so optimal solutions) can be achieved with smaller population sizes, leading to a significant decrease in MOGA's runtime. To this end, we implemented an algorithm based on the Pareto Front method to select the configuration that strikes a balance between the 'fitness value of the best solution' and the 'MOGA's runtime'.

To start the algorithm, MOGA is run for every combination of 'population-size', 'elitism-rate', 'cross-over rate', and 'mutation rate' for a very limited number of scenarios. After each run, (1) the fitness value of the best solution and (2) MOGA's runtime are recorded for further processing. Figure 4 shows sample records after collecting all points for different sizes of problem instances. Using these points, we aim to find several formulas to relate each MOGA's parameter to the characteristics of a problem instance; for example, finding a formula to select a reasonable population size for MOGA based on the number of user-helpers and total number of computing nodes in an infrastructure. Using such formulas to set MOGA's input parameters, we can then ensure that MOGA will (most likely) find the best solution (or one very close to it) within a reasonable amount of time.

To find a point that strikes a balance between the solution quality and MOGA's runtime, we first find the Pareto Front (curve) of these points. To that end, a configuration (out of all available points, shown in blue in Figure 4) is randomly selected, and all other

configurations whose fitness and runtime are worse than the selected point are deleted. This process is repeated until no further points can be removed from the list. The calculated Pareto-front that represent 'best configurations' (because no point is strictly better than any other point on the same Pareto front) are then sorted based on their population size, and their median is identified as the optimal population size for that scenario. Similarly, the same 'best configurations' are then re-sorted based on their crossover rate, and their median is identified as the optimal crossover rate for that scenario. The same procedure (re-sorting) is repeated to select the best elitism and mutation rates.

Figure 4 illustrates how our Pareto Front-based approach selects the best configuration to balance fitness and run-time on various scales. Based on the results in Figure 4, we observe that although the fitness value of the best configuration points in terms of fitness are relatively superior compared to the best configuration points selected by the Pareto Front-based method, these superiorities are negligible compared to the variations in runtime. Specifically, the optimal configuration point selected by the Pareto Front-based method shows 63%, 57%, 60%, and 53% better runtime in small-scale, medium-scale, large-scale and xLarge-scale scenarios, respectively. On the contrary, the best fitness configuration points only show improvements of 6%, 4%, 3%, and 7% compared to Pareto Front-based configuration points on all scales, respectively.

The optimal configurations considered by our Pareto Front-based method for all scenarios are indicated in Table 10. As shown in Table 10, the optimal

Table 10: Optimal values for MOGA parameters.

| Parameter | Small-scale | Medium-scale | Large-scale | xLarge-scale |
|-----------|-------------|--------------|-------------|--------------|
| <i>ps</i> | 200 | 300 | 400 | 500 |
| <i>cr</i> | 60% | 70% | 70% | 80% |
| <i>mr</i> | 1% | 1% | 1% | 1% |
| <i>ss</i> | 20 | 30 | 40 | 50 |
| <i>it</i> | 50 | 100 | 150 | 200 |

population size increases slightly with the problem size. On the contrary, the optimal value of the mutation rate and the selection size are considered 1% and 10% for all scenarios, respectively. Furthermore, the crossover rate falls within the range 60% to 80%, which means that a larger problem size requires a relatively higher crossover rate. Using the information in Table 10, we can provide an equation to estimate the optimal value for population size. Equation (22) estimates the optimal population size in various problem sizes, where $(X \times Y)$ is the chromosome length, where N , M , and K represent the total number of user nodes, helper nodes, and computing nodes, respectively. α , β , and γ are fixed values that we calculated (using the curve fitting algorithm) $\alpha = 0.9$, $\beta = 0.16$ and $\gamma = 0.16$ to estimate the optimal population size.

$$ps \approx 100^\alpha \cdot (X \times Y)^\beta \cdot (N + M + K)^\gamma \quad (22)$$

Similarly, the crossover rate (*cr*) can be estimated using Equation (23), where c is calculated using Equation (24) where $\delta = 0.0003$ and $\varepsilon = 0.04$. Since we identified an optimal crossover rate 60% to 80% through our fine-tuning process for MOGA, in Equation (23), we consider specific conditions to set the crossover rate within the range 60% to 80%.

$$cr \approx \begin{cases} c & c \leq 0.8 \\ 0.8 & c > 0.8 \end{cases} \quad (23)$$

$$c = 0.6 + (X \times Y) \cdot \delta + (N + M + K)^\varepsilon \quad (24)$$

Figure 5 shows the MOGA convergence process in terms of the best, median, and worst chromosomes of the population in the different scenarios. It is evident that MOGA converges rapidly on the small and medium scale, particularly within the first 100 iterations. Therefore, we set the number of MOGA iterations to 50 and 100 for small and medium scales to achieve better run-time efficiency. We also set the number of MOGA iterations at 150 and 200 for the large and xLarge scales, respectively, because no notable improvements in the fitness value were observed across these scenarios after 150 and 200 iterations.

As a result, as the problem size increases, MOGA requires more iterations. We also developed Equation (25) to estimate sufficient number of iterations for MOGA, where $\zeta = 0.9$, $\eta = 0.69$, and $\theta = 0.1$.

$$it \approx 50^\zeta + (X \times Y)^\eta \cdot (N + M + K)^\theta \quad (25)$$

Figure 5 provides more information on the overall performance of the MOGA. We investigated the median and worst chromosomes to assess the performance of MOGA in terms of exploration and exploitation. The results show that the median chromosomes consistently follow the best chromosomes, indicating that MOGA effectively exploits good solutions. The worst chromosomes also show an improvement over time, indicating that the algorithm maintains diversity, while still exploring different regions of the solution space.

6.4 Response Time Evaluation

Figure 6 (a) compares the total response time achieved by MOGA with the heuristic algorithms in our four scenarios. More specifically, MOGA leads to an average improvement of 71%, 67%, 65%, and 66% in total response time across small-scale, medium-scale, large-scale, and xLarge-scale, respectively. Furthermore, Figure 6 (b) compares the average response time of the algorithms for each service component. As depicted in Figure 6 (b), the average response time of MOGA for a service component is approximately 2 seconds on the different scales, while the heuristic algorithms indicate average response times above 5 seconds. The obtained results related to response time show the superior performance of our MOGA compared to the heuristic algorithms and indicate MOGA's capability (as a meta-heuristic algorithm) to identify and implement efficient solutions with minimal response times.

6.5 Infrastructure and Service Reliability Evaluation

Figure 6 (c) illustrates a comparison of the infrastructure reliability achieved by both MOGA and the heuristic algorithms. On average, MOGA shows an infrastructure reliability close to 98%, whereas the reliability achieved by other heuristic algorithms falls below 91%. Furthermore, Figure 6 (d) compares the service reliability between MOGA and the heuristic algorithms. The results in Figure 6 (d) indicate that MOGA consistently maintains a high level of service reliability in all scenarios, with an average of 96%. In addition, a high level of service reliability is observed for MR in all scenarios, with an average of 97%. As

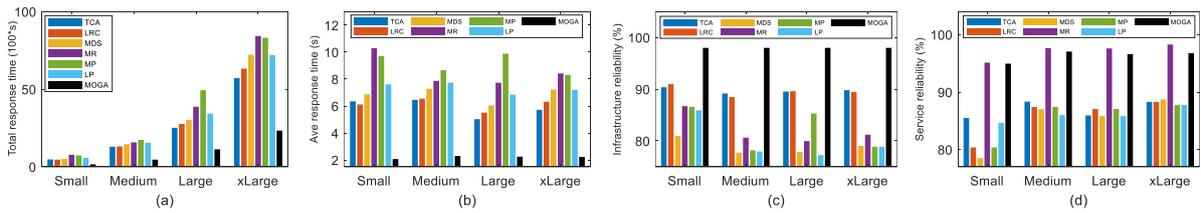


Figure 6: Evaluations based on response time and system reliability.

Table 11: Runtime of algorithms (s).

| Algorithm | Small-scale | Medium-scale | Large-scale | xLarge-scale |
|-----------|-------------|--------------|-------------|--------------|
| TCA | 0.0026 | 0.0109 | 0.0411 | 0.0499 |
| LRC | 0.0014 | 0.0026 | 0.0087 | 0.0239 |
| MDS | 0.0021 | 0.0317 | 0.0464 | 0.0622 |
| MR | 0.0015 | 0.0328 | 0.0299 | 0.0899 |
| MP | 0.0024 | 0.0058 | 0.0301 | 0.0951 |
| LP | 0.0013 | 0.0351 | 0.084 | 0.1075 |
| MOGA | 8.73 | 79.74 | 530.78 | 2678.25 |

mentioned above, MR prioritizes the placement of service components with higher reliability scores, resulting in high service reliability results. However, MR faces limitations in providing desirable infrastructure reliability and response time. Regarding the other heuristic algorithms, they provide a service reliability ranging from 75% to 88%.

Based on the findings presented in this subsection, MOGA provides a high degree of infrastructure reliability and an excellent degree of service reliability. This maximization in reliability, both in hardware and software aspects, is vital in minimizing interruptions arising from reliability-related issues within edge-to-cloud environments, mainly when it involves the demanding requirements of AR/VR applications. The results clearly illustrate the scalability of MOGA that consistently maintains a high degree of infrastructure and service reliability on various scales of the system. It indicates its adaptability and effectiveness regardless of the scale of the system.

6.6 Algorithm Runtime Evaluation

Table 11 presents the runtime of both MOGA and the heuristic algorithms. As expected, the run-time of MOGA is higher compared to the heuristic algorithms due to the nature of its operators. However, we made a significant reduction in MOGA runtime by optimizing its configurations.

6.7 Distribution of Service Components over the Infrastructure

Figure 7 depicts the percentage of distribution of service components using the algorithms and provides detailed information on the operation of the algorithms.

Based on the results of Figure 7, it is clear that MOGA never assigns service components for execution on helper nodes. The reason is that the helper nodes communicate with the users through the cloud layer. Therefore, it is rational to place service components on superior nodes in proximity to user nodes rather than helper nodes. Additionally, MOGA places only a minor portion of the service components to user nodes. In fact, MOGA minimizes the usage of user nodes due to their constrained computational and storage capacities. Instead, MOGA prefers to employ more robust computing nodes in Tier-1 and Tier-2, which are located near user nodes, to improve performance. Therefore, we can conclude that MOGA minimizes response time by decreasing the placement of service components in Tier-3 and the rare usage of helper and user nodes. This not only leads to a decrease in the overall transmission time but also reduces execution time with respect to leveraging more powerful computing nodes at the network's edge. However, it is essential to note that MOGA also optimally selects the appropriate version of the service component on the appropriate computing nodes in each tier, leading to significant improvements both in response time and system reliability. Although the selection of nodes in Tiers 2 and 3 seems trivial, and thus may falsely give the belief that the cost of running MOGA (or other meta-heuristics for that matter) is not worth, it is essential to note that other heuristics that mentioned in this paper and roughly deploy such simplistic procedures could not lead to MOGA's high quality answers. This is mainly due to the fact that simultaneously selecting both computing nodes and software versions adds extra layers of complexity that simple heuristics cannot tolerate.

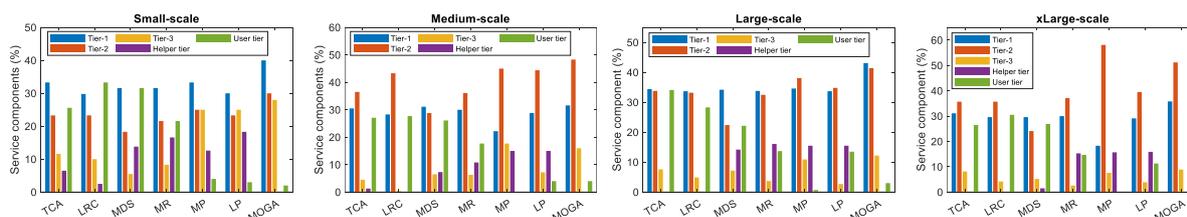


Figure 7: Distribution of service components over the infrastructure in different scales.

7 CONCLUSION

In this paper, we introduce a MOGA for optimal service placement in edge-to-cloud AR/VR systems. The primary objectives were to minimize service response time, maximize infrastructure reliability, and achieve the highest service reliability by optimally placing service components on computing nodes, user nodes, and helper nodes in the edge-to-cloud infrastructure. We devised a robust fine-tuning strategy to attain optimal configurations for our MOGA in order to strike a balance between the MOGA's runtime and the quality of its solutions. We also implemented a simulator to validate the proposed MOGA's effectiveness. Through extensive simulations and measurements on various scales, we showed the importance of designing MOGA (as well as other meta-heuristic) to simultaneously optimize for response time and infrastructure and service reliability. We also assessed the performance of MOGA in terms of the distribution of service components over the infrastructure, illustrating how our MOGA provides an optimal solution for the placement of AR/VR services compared to the other rather simple heuristic algorithms.

ACKNOWLEDGEMENTS

Parts of this work have been supported by the Knowledge Foundation of Sweden (KKS).

REFERENCES

- Abedi, S., Ghobaei-Arani, M., Khorami, E., and Mojarad, M. (2022). Dynamic resource allocation using improved firefly optimization algorithm in cloud environment. *Applied Artificial Intelligence*, 36(1):2055394.
- Acheampong, A., Zhang, Y., and Xu, X. (2023). A parallel computing based model for online binary computation offloading in mobile edge computing. *Computer Communications*, 203:248–261.
- Amini Motlagh, A., Movaghar, A., and Rahmani, A. M. (2022). A new reliability-based task scheduling algorithm in cloud computing. *International Journal of Communication Systems*, 35(3):e5022.
- Chen, X., Xu, H., Zhang, G., Chen, Y., and Li, R. (2022). Unsupervised deep learning for binary offloading in mobile edge computation network. *Wireless Personal Communications*, 124(2):1841–1860.
- Cozzolino, V., Tonetto, L., Mohan, N., Ding, A. Y., and Ott, J. (2022). Nimbus: Towards latency-energy efficient task offloading for ar services. *IEEE Transactions on Cloud Computing*.
- De Souza, A. B., Rego, P. A. L., Chamola, V., Carneiro, T., Rocha, P. H. G., and de Souza, J. N. (2023). A bee colony-based algorithm for task offloading in vehicular edge computing. *IEEE Systems Journal*.
- Dong, L., Wu, W., Guo, Q., Satpute, M. N., Znati, T., and Du, D. Z. (2019). Reliability-aware offloading and allocation in multilevel edge computing system. *IEEE Transactions on Reliability*, 70(1):200–211.
- Elawady, M. and Sarhan, A. (2020). Mixed reality applications powered by ioe and edge computing: A survey. In *Internet of Things—Applications and Future: Proceedings of ITAF 2019*, pages 125–138. Springer.
- Fan, W., Zhao, L., Liu, X., Su, Y., Li, S., Wu, F., and Liu, Y. (2022). Collaborative service placement, task scheduling, and resource allocation for task offloading with edge-cloud cooperation. *IEEE Transactions on Mobile Computing*.
- Fang, D., Xu, H., Yang, X., and Bian, M. (2020). An augmented reality-based method for remote collaborative real-time assistance: from a system perspective. *Mobile Networks and Applications*, 25:412–425.
- Herabad, M. G. (2024). service-placement-simulator. <https://github.com/ms-garshasbi/service-placement-simulator>.
- Huang, Z. and Friderikos, V. (2021). Proactive edge cloud optimization for mobile augmented reality applications. In *2021 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1–6. IEEE.
- Huangpeng, Q. and Yahya, R. O. (2024). Distributed iot services placement in fog environment using optimization-based evolutionary approaches. *Expert Systems with Applications*, 237:121501.
- Ji, T., Wan, X., Guan, X., Zhu, A., and Ye, F. (2023). Towards optimal application offloading in heterogeneous edge-cloud computing. *IEEE Transactions on Computers*.
- Khaleel, M. I. (2022). Multi-objective optimization for scientific workflow scheduling based on performance-to-

- power ratio in fog–cloud environments. *Simulation Modelling Practice and Theory*, 119:102589.
- Lavanya, M., Shanthi, B., and Saravanan, S. (2020). Multi objective task scheduling algorithm based on sla and processing time suitable for cloud environment. *Computer Communications*, 151:183–195.
- Li, J., Shang, Y., Qin, M., Yang, Q., Cheng, N., Gao, W., and Kwak, K. S. (2022). Multiobjective oriented task scheduling in heterogeneous mobile edge computing networks. *IEEE Transactions on Vehicular Technology*, 71(8):8955–8966.
- Liu, G., Zeng, Y., Li, D., and Chen, Y. (2015). Schedule length and reliability-oriented multi-objective scheduling for distributed computing. *Soft Computing*, 19:1727–1737.
- Liu, J., Yang, P., and Chen, C. (2023a). Intelligent energy-efficient scheduling with ant colony techniques for heterogeneous edge computing. *Journal of Parallel and Distributed Computing*, 172:84–96.
- Liu, Y., Mao, Y., Liu, Z., Ye, F., and Yang, Y. (2023b). Joint task offloading and resource allocation in heterogeneous edge environments. *IEEE Transactions on Mobile Computing*.
- Maciel, P., Dantas, J., Melo, C., Pereira, P., Oliveira, F., Araujo, J., and Matos, R. (2021). A survey on reliability and availability modeling of edge, fog, and cloud computing. *Journal of Reliable Intelligent Environments*, pages 1–19.
- Madni, S. H. H., Latiff, M. S. A., Ali, J., and Abdulhamid, S. M. (2019). Multi-objective-oriented cuckoo search optimization-based resource scheduling algorithm for clouds. *Arabian Journal for Science and Engineering*, 44:3585–3602.
- Mahjoubi, A., Grinnemo, K.-J., and Taheri, J. (2022a). An efficient simulated annealing-based task scheduling technique for task offloading in a mobile edge architecture. In *2022 IEEE 11th International Conference on Cloud Networking (CloudNet)*, pages 159–167. IEEE.
- Mahjoubi, A., Grinnemo, K.-J., and Taheri, J. (2022b). Ehga: A genetic algorithm based approach for scheduling tasks on distributed edge-cloud infrastructures. In *2022 13th International Conference on Network of the Future (NoF)*, pages 1–5. IEEE.
- Moghaddas, V., Fazeli, M., and Patooghy, A. (2016). Reliability-oriented scheduling for static-priority real-time tasks in standby-sparing systems. *Microprocessors and Microsystems*, 45:208–215.
- Raju, M. R. and Mothku, S. K. (2023). Delay and energy aware task scheduling mechanism for fog-enabled iot applications: A reinforcement learning approach. *Computer Networks*, 224:109603.
- Ren, J., Gao, L., Wang, X., Ma, M., Qiu, G., Wang, H., Zheng, J., and Wang, Z. (2021). Adaptive computation offloading for mobile augmented reality. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(4):1–30.
- Siriwardhana, Y., Porambage, P., Liyanage, M., and Ylianttila, M. (2021). A survey on mobile augmented reality with 5g mobile edge computing: Architectures, applications, and technical aspects. *IEEE Communications Surveys & Tutorials*, 23(2):1160–1192.
- Wang, L., Deng, X., Gui, J., Chen, X., and Wan, S. (2023). Microservice-oriented service placement for mobile edge computing in sustainable internet of vehicles. *IEEE Transactions on Intelligent Transportation Systems*.
- Yeganeh, S., Sangar, A. B., and Azizi, S. (2023). A novel q-learning-based hybrid algorithm for the optimal offloading and scheduling in mobile edge computing environments. *Journal of Network and Computer Applications*, 214:103617.